

# Magellan Cloud at ALCF

Susan Coghlan

Argonne Leadership Computing Facility



@



# Magellan Background

## DOE Midrange Computing Report

“Midrange computing, and the associated data management play a vital and growing role in advancing science in disciplines where capacity is as important as capability.”

“Demand seems to be limited only by the availability of computational resources.”

“The number of alternative ways for providing these capabilities is increasing.”



***From: Mid-range Computing in the Support of Science at Office of Science Laboratories. Report of a Workshop, October 2008***



## Some Midrange Computing Sweet Spots

- Serial or scalability-challenged codes
- Science that does not require tight coupling
  - Trivially parallel app, Parameter sweeps, Monte Carlo methods
- Science that can run at low-concurrency
  - 2D v. 3D, different scales for different steps, parameter validation
- On-ramp to the large centers
  - Training, code development, staging
- Data-intensive science
  - Includes Real-time, Visualization



## Some Midrange Computing Sweet Spots

- Serial or scalability-challenged codes
- Science that does not require tight coupling
  - Trivially parallel app, Parameter sweeps, Monte Carlo methods
- Science that can run at low-concurrency
  - 2D v. 3D, different scales for different steps, parameter validation
- On-ramp to the large centers
  - Training, code development, staging
- Data-intensive science
  - Includes Real-time, Visualization

Could cloud computing be used to address these needs?





# What is Cloud Computing?

- “Infinite” number of nodes
- On demand access
- No locality
- Virtual machines
  - Your personal software stack
- Nodes are independent resources



Amazon EC2 is the canonical example



# The Dark Side of Clouds

- Interconnect suitable only for loosely coupled applications
- Non-uniform execution times (VM jitter)
- Practical limits to the size of a cluster
- Poor shared disk I/O
- Substantial data storage and I/O costs
- Still self-supported



**These issues are not intrinsic to clouds, only current implementations.**



# Magellan Cloud Project

- Funded by DOE under the American Recovery and Reinvestment Act (ARRA)
- Joint project between Argonne Leadership Computing Facility (ALCF) and National Energy Research Scientific Computing Center (NERSC)
- Funds for hardware testbeds located at both NERSC and Argonne
- 2 year research project to investigate whether clouds could be used to address DOE Midrange computing needs



# Magellan Goal: To Explore!

- Discover science applications and user communities well suited for cloud computing
- Understand the deployment and support issues required to build large science clouds
  - Is it cost effective and practical to operate science clouds?
  - How can commercial clouds be leveraged? (stretch)
- Pioneer new cloud software and infrastructure that can better meet the needs of science
- Investigate how software as a service from the clouds can support science
- Examine how cloud computing can support data-intensive science
- Explore the challenges for security in a virtualized cloud environment



# Unique Characteristics of Magellan

- High Speed, Low Latency Interconnect
  - QDR Infiniband Connection to All Nodes
- High Performance Storage
  - Solid State Storage
  - High Performance Parallel Filesystem
- High Bandwidth Wide Area Networking
  - Direct Connection to 100Gbps ANI
- Tuned Middleware and Scientific Software



# Argonne Magellan Cloud Hardware - Phase 1


## Spring 2010

### Compute

504 Compute Nodes  
Nehalem Dual quad-core 2.66GHz  
24GB RAM, 500GB Disk  
QDR IB link

### Totals

4032 Cores, 40TF Peak  
12TB RAM, 250TB Disk



# Argonne Magellan Cloud Hardware - Phase 2 Late Summer 2010

## Compute

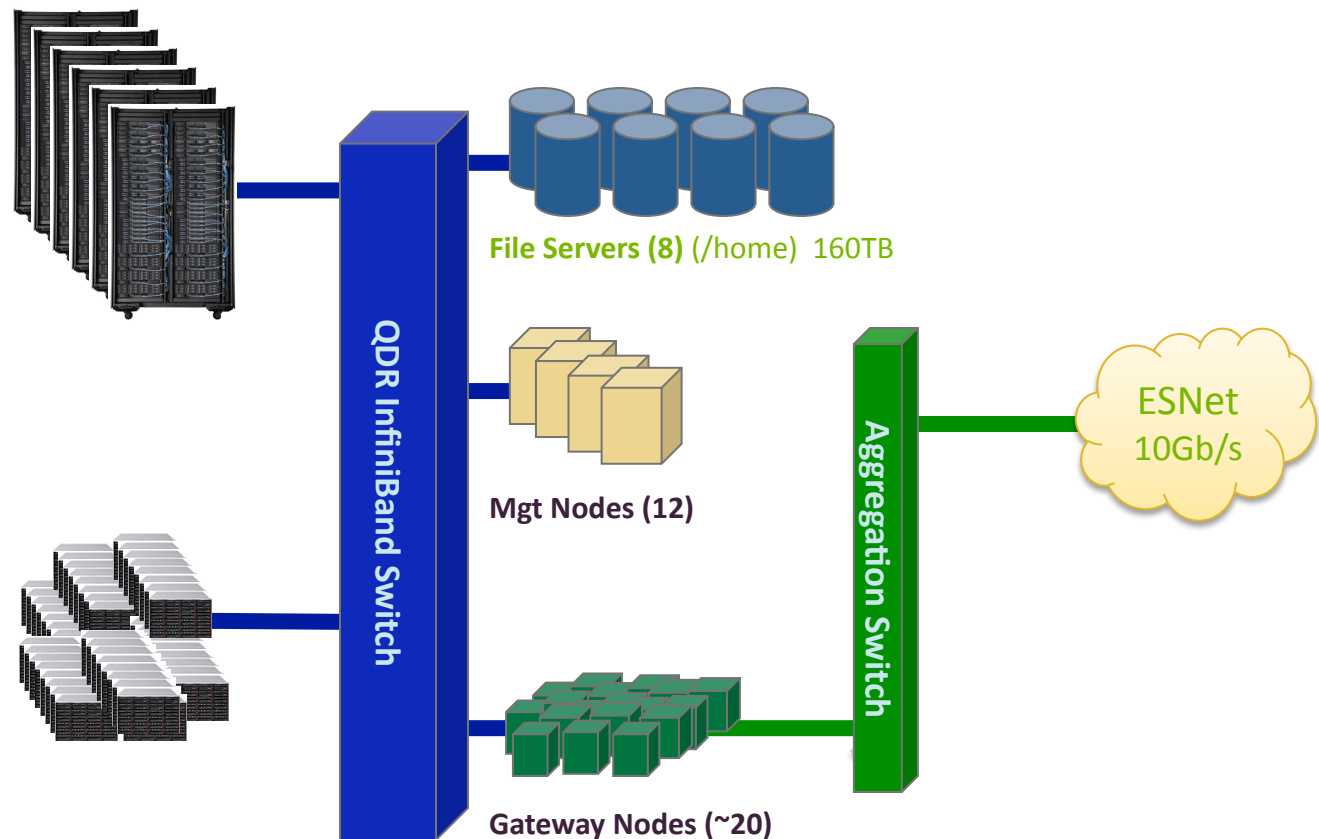
504 Compute Nodes  
Nehalem Dual quad-core 2.66GHz  
24GB RAM, 500GB Disk  
QDR IB link

## Totals

4032 Cores, 40TF Peak  
12TB RAM, 250TB Disk

## Active Storage

~100 Compute/Storage Nodes  
~10TB FLASH/SSD Storage  
~500TB Disk Storage



# Argonne Magellan Cloud Hardware - Final

## January 2011

### Compute


504 Compute Nodes  
Nehalem Dual quad-core 2.66GHz  
24GB RAM, 500GB Disk  
QDR IB link

### Totals

4032 Cores, 40TF Peak  
12TB RAM, 250TB Disk

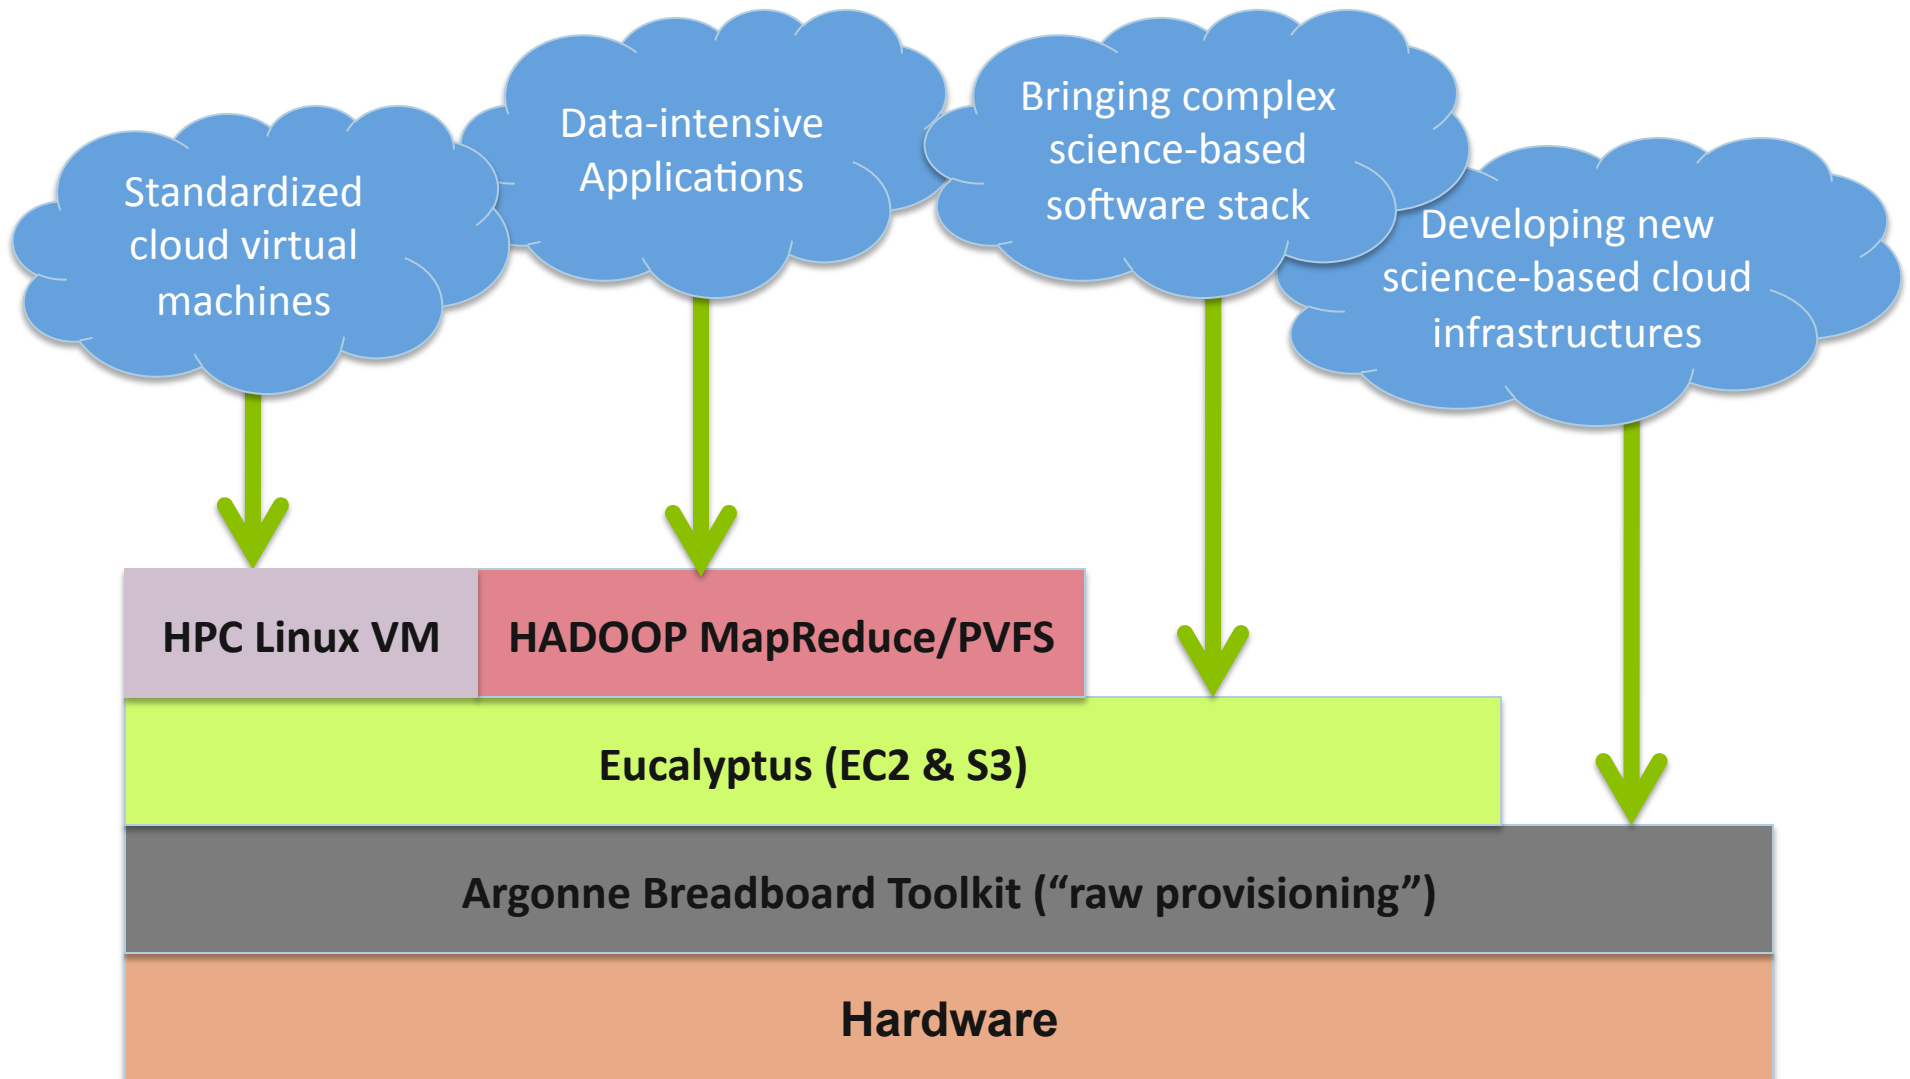
### Active Storage

~100 Compute/Storage Nodes  
~10TB FLASH/SSD Storage  
~500TB Disk Storage





# Argonne Magellan Software Architecture



# Magellan Users

- Science teams
  - Explore using the cloud to accomplish science
- ANI 100Gb/s project teams
  - Add unique capabilities to cloud computing with fast networking
- Application Developers
  - Develop applications and ensemble pipelines well-suited for clouds
- System software developers
  - Develop software that provides unique capabilities for science applications running in the cloud
- System manager
  - Integrate transient resources into local production offering
- System operations
  - Provide secure environment and coordinate ALCF/NERSC services



# ALCF Magellan High-Level Time Line

- January 2010 – Core System Delivered
- March 2010 – Open to Early Users, Welcome to Magellan Day (3/23)
- April 2010 – Open to regular users
- September 2010 – Phase 2 hardware delivered, Welcome to Magellan Day II
- November 2010 – Phase 2 hardware open to regular users
- January 2011 – 100Gbit deployed
- October 2011 – Project ends



# We're Looking for a Few Good Users

- Projects That
  - Exercise Unique Properties of Magellan and Clouds
  - Leverage Research at ALCF and MCS
  - Are Not Well-Served by Current Resources
- People That
  - Tolerate Downtime and Configuration Changes
  - Provide Feedback and Cooperative Experimentation
  - Have the Time to be Part of the Magellan Team



■ <http://magellan.alcf.anl.gov>



**Magellan**  
a cloud for science

Home Architecture Science Cloud Research User Support

Active Storage  
~100 Compute/Storage Nodes; ~10TB FLASH/SSD Storage; ~500TB Disk Storage

Scientists Look to the Clouds  
Benefiting from mid-range computing resources and data-storage systems

[Request an Allocation](#)

**About Magellan**

Magellan is a research and development effort to establish a nationwide scientific mid-range distributed computing and data analysis testbed. It has two sites (NERSC and ALCF) with multiple 10's of teraflops and multiple petabytes of storage, as well as appropriate cloud software tuned for moderate concurrency.

Results from both sites will be compared to existing mid-range resources in the Office of Science labs. These results will generate data for a cost-benefit analysis of various mid-range computing options for the Office of Science.

**Our Goals**

Through this project we hope to promote open interface specifications for clouds, as well as to determine:

1. turnaround time for results.
2. if the Magellan cloud produces science results as fast or faster than local clusters.
3. ease of use.
4. if the interface allows the job to be directed to different resources or split between resources without requiring additional input from the user.

**News**

- ① U.S. government explores limits of cloud and virtualization
- ② DOE to explore scientific cloud computing at Argonne, Lawrence Berkeley national laboratories.
- ③ National Impact Series: Scientists Look to the Clouds to Solve Complex Questions

WMD

Welcome to  
**Magellan Day**  
Tuesday, March 23rd, 2010  
**9:00 am - 3:45 pm**  
ICS Building 240 Conference Center  
Room 1406/1407

